

THE BELL SYSTEM TECHNICAL JOURNAL

DEVOTED TO THE SCIENTIFIC AND ENGINEERING
ASPECTS OF ELECTRICAL COMMUNICATION

Volume 59

February 1980

Number 2

Copyright © 1980 American Telephone and Telegraph Company. Printed in U.S.A.

A Single-Chip VLSI Echo Canceler

By D. L. DUTTWEILER and Y. S. CHEN

(Manuscript received August 2, 1979)

A single-chip VLSI (very large-scale integration) echo canceler has been fabricated in NMOS (N-channel metal-oxide semiconductor). The canceler has a 128-tap (16-ms) delay line and a white-noise convergence rate of 70 dB/s. The chip measures 313 by 356 mils and contains 35,000 devices.

I. INTRODUCTION

On long-distance telephone calls, echo is an impairment that is often as annoying subjectively as the more obvious impairments of low volume and noise. Since the subjective annoyance of echo increases with delay as well as level, the measures used for its control depend on circuit length.

For circuits under about 2000 miles, via net loss (VNL) is used.¹ The via-net-loss plan inserts increasing amounts of loss with increasing circuit length. Since the loss is inserted in the transmission paths for both directions, echo is attenuated by $2 \times \text{VNL}$, whereas the desired signal is only attenuated by $1 \times \text{VNL}$. Thus, the signal-to-echo ratio improves by $1 \times \text{VNL}$.

For circuits over about 2000 miles, the amount of VNL needed to give satisfactory echo performance would result in unacceptable received levels.² On terrestrial circuits over this length, echo suppressors are used to control echo. An echo suppressor is basically a voice-operated switch that attempts to open the transmission path from the listening customer to the talking customer. Two problems associated with echo suppressors are their inability to control echo when both customers speak simultaneously (double-talk) and the chopping and

clipping they impart to weak speech. In spite of these drawbacks, echo suppressors generally are judged to provide acceptable echo control on long terrestrial circuits.

Because of the high altitude (22,000 miles) of geostationary satellites, satellite circuits introduce extremely long echo delays. On a full-hop satellite circuit, that is, a circuit carrying both directions of transmission via satellite, the echo delay is 540 ms. This delay can be halved by using satellite circuits half-hop with one direction of transmission on terrestrial facilities. Extensive subjective testing³ has shown that whereas echo suppressors still provide acceptable echo control on half-hop circuits, their performance on full-hop circuits is unacceptable for domestic service. For this reason, domestic satellite circuits are currently deployed half-hop. The half-hop solution is not entirely satisfactory, however. It is difficult to administer, and it does not allow as extensive a use of satellite circuits as full-hop would allow.

Echo cancelers, which control echo by creating an echo replica and subtracting it from the return signal (see Section II), offer a better solution. Since the initial work^{4,5} on echo cancelers around 1965, their potential for echo control on full-hop satellite circuits has been recognized and, indeed, the tests referred to above³ showed full-hop satellite circuits equipped with echo cancelers to be almost equal in subjective quality to long-distance terrestrial circuits. The problem has been cost. The echo canceler to be described requires 250 words of storage, a precision multiplier and accumulator capable of one million operations per second, and extensive miscellaneous control logic (which, interestingly, dominates the chip area). A number of echo cancelers were realized around 1970,⁶⁻¹² but these cancelers were built for laboratory experimentation rather than production. Cancelers seriously proposed for manufacture were first described in the technical literature about 2 years ago,¹³⁻¹⁶ but even these cancelers had projected or actual per-channel costs that were high enough to discourage wide-spread application. Now that an integrated echo canceler has been realized and the cost of echo canceling has become almost inconsequential for many applications, widespread deployment of echo cancelers on satellite circuits, and perhaps even nonsatellite circuits as well, is likely.

Following this introductory section, there are four other sections entitled

- II. Echo Canceler Basics
- III. Chip Characteristics
- IV. Performance Data
- V. Acknowledgments.

II. ECHO CANCELER BASICS

Figure 1 shows a satellite circuit with no echo protection. Hybrids at both ends of the circuit convert from the 2-wire transmission used on




Fig. 1—Satellite circuit with no echo protection.

customer loops and metallic trunks to the 4-wire transmission needed for carrier circuits. The 270-ms delay in each 4-wire path represents the satellite circuit delay. Ideally, when the customer on the left speaks, his speech should follow the upper transmission path to the right-hand hybrid and from there be directed to the 2-wire circuit. In actuality, however, not all the speech energy will be directed to the 2-wire circuit and some will be returned along the lower 4-wire path to be heard as an echo delayed by 540 ms.

Echo cancelers are installed in the network in pairs in exactly the same manner as split echo suppressors (see Fig. 2). Figure 3 shows a single canceler with labels for important signals. The underlying assumption in echo canceling is that the echo return path, from the point where the canceler bridges to the point where the echo estimate is subtracted, is linear and time-invariant. Under this assumption, the return signal $y(k)$ is expressible as

$$y(k) = h_k * x(k) + v(k), \quad (1)$$

where $x(k)$ is the far-end speech, $v(k)$ is the near-end speech plus any additive noise, h_k is the impulse response of the echo path, and $*$ denotes convolution. The echo canceler forms an estimate \hat{h}_k of the echo path impulse response and then estimates the echo as

$$\hat{y}(k) = \hat{h}_k * x(k). \quad (2)$$

The internal structure of an echo canceler is shown in Fig. 4. A tapped delay line stores the N most recent samples of $x(k)$, and the signal at each position along this delay line is multiplied by a tap weight $\hat{h}_n(k)$. The sum

$$\hat{y}(k) = \sum_{n=0}^{N-1} \hat{h}_n(k) x(k-n) \quad (3)$$

is the echo estimate.

Since the true echo-path impulse response is necessarily causal, we have, from (1),

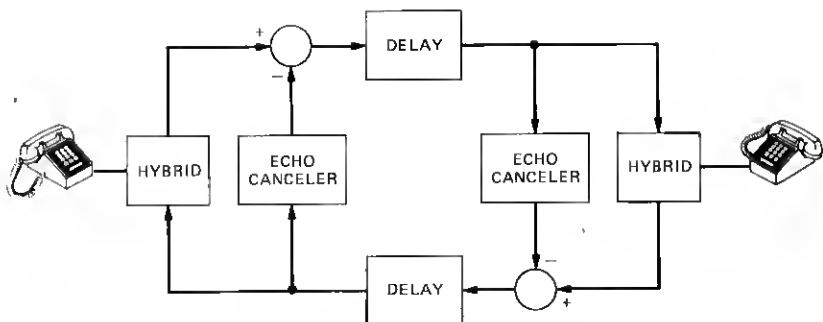


Fig. 2—Using echo cancelers to control echo.




Fig. 3—Signal definitions.

$$y(k) = \sum_{n=0}^{\infty} h_n x(k-n) + v(k). \quad (4)$$

Thus, if the echo-path impulse response decays to zero after N samples and

$$\hat{h}_n(k) = h_n, \quad n = 0, 1, \dots, N-1, \quad (5)$$

the echo in $y(k)$ will be canceled.

Feedback circuitry, shown in Fig. 4 for the first tap position only, drives $\hat{h}_n(k)$ toward h_n at each tap position n . Each sample period, the tap weights are updated as

$$\hat{h}_n(k+1) = \hat{h}_n(k) + Kx(k-n)e(k), \quad (6)$$

where K is the loop gain and

$$e(k) = y(k) - \hat{y}(k) \quad (7)$$

is the error signal. Proofs that this algorithm does converge the tap weights as desired appear in Refs. 17 and 18 and many other papers on adaptive filtering.

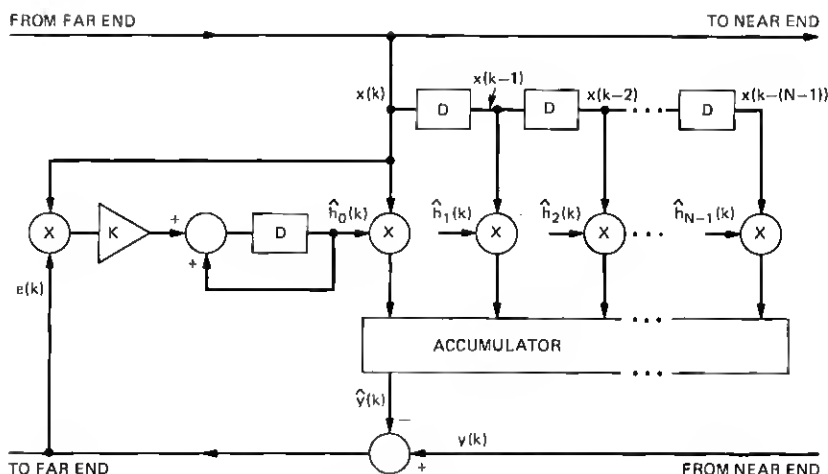


Fig. 4—Echo canceler structure.

A typical impulse response h_k is shown in Fig. 5. The time span over which the response differs significantly from zero is generally about 2 ms (16 taps with 8-kHz sampling). Because of flat delay in the transmission path from the echo canceler to the hybrid, the nonzero region for the impulse response does not begin at time zero, but is delayed. The amount of end delay that must be allowed for in the "tail" of the circuit dictates N . With $N = 128$, as in the VLSI echo canceler, tails of up to 1000 miles can be accommodated.

Another key parameter is the loop gain K . If K is made larger, the updates to the tap weights will be larger and convergence will be quicker. The disadvantage in using larger loop gains is that the estimates $\hat{h}_n(k)$ are then noisier and the echo cancellation is not as good. In fact, if K is made too large, instability results.

There is approximately a 30-dB dynamic range between weak and strong talkers in the telephone plant. In Ref. 16, it is argued that, because of this large dynamic range, it is desirable to normalize the loop gain K as

$$K = K' / \hat{\sigma}^2(k), \quad (8)$$

where K' is a normalized loop gain and $\hat{\sigma}^2(k)$ is an estimate of the variance (power) of $x(k)$. The VLSI echo canceler has loop-gain normalization with K' approximately equal to 0.001. With white noise at the $x(k)$ input, we have from equation (15) of Ref. 16.

$$\begin{aligned} \tau &= 1/K' \text{ sample periods} \\ &= 1000 \text{ sample periods} \\ &= 0.125 \text{ second,} \end{aligned} \quad (9)$$




Fig. 5—Typical impulse response.

where τ is the time constant of adaptation. Thus the rate of convergence with a white noise input is

$$20 \log e = 8.7 \text{ dB} \quad (10)$$

in $\frac{1}{8}$ second, or approximately 70 dB/s.

A normalized loop gain of $K' = 0.001$ is comparatively small. We could have chosen a much larger loop gain without making the impulse response estimates excessively noisy (see equation (16) of Ref. 16). A comparatively small value was chosen, since small K' eases the requirements on near-end speech detection (see the following paragraph) and there does not appear to be any subjective advantage in convergence times faster than a few syllables of speech. (The VLSI echo canceler is not intended for use where phase-roll¹⁹ might be encountered.) It is interesting to note that, except for the problem of realizing a satisfactory near-end speech detector, it is easier to realize a high-gain echo canceler than a low-gain echo canceler. The reason is that a low-gain echo canceler, which averages updates over a longer time, requires additional lower-order bits in the digital representations of the tap weights \hat{h}_n .

Near-end speech $v(k)$ in the return signal $y(k)$ will typically be even stronger than the echo $h_k * x(k)$. This near-end speech is unwanted noise as far as the convergence algorithm is concerned. It would diverge the echo canceler if the canceler were to continue updating tap weights while it was present. A critical component of an echo canceler is thus a near-end speech detector to inhibit updates while near-end speech is present. It is important to realize that, unlike an echo suppressor, which cannot control echo during double-talk, an

echo canceler is still canceling echo during double-talk. Near-end speech is troublesome to an echo canceler only in that the canceler is unable to improve the current echo path estimate while it is present.

III. CHIP CHARACTERISTICS

A photomicrograph of the canceler chip appears in Fig. 6. The dense block at the left center consists of some 3000 bits of serial shift register, while the other area contains approximately 3000 gates of random logic. The total transistor count is close to 35,000. The chip measures 313 by 356 mils and fits in a 24-pin hermetic dip. It is realized using an




Fig. 6—Photomicrograph of echo canceler.

enhancement NMOS technology that involves five-mask levels and 5 micron line widths. Two power supplies are used: V_{DD} at $5V \pm 10$ percent and V_{GG} at $12V + 10$ percent and -14 percent. The substrate is grounded. The operating ambient temperature extremes are 0°C and 85°C . All I/Os are LSTTL (low-power Schottky transistor-transistor logic)-compatible.

The random logic portion of the chip was laid out using a modular approach. Logic elements such as NORs and flip-flops had predefined lay-outs and were precharacterized for logic, timing, and fault-coverage simulators. The steps required to realize the echo canceler chip were to

- (i) Design the logic using the available modular building blocks.
- (ii) Verify the logic using the logic and timing simulators.
- (iii) Define test-vector sequences.
- (iv) Use the fault-coverage simulator to verify that the test-vector sequences provided adequate fault coverage.
- (v) Position the modular building blocks on the chip in such a way that the chip area required for interconnect wiring was close to minimum.

This modular approach can lead to successful chips on the first try, and such was the case for the echo canceler.

The shift register memory is realized with a custom two-phase ratioless dynamic circuit. The shift registers are laid out in an array with power busses running perpendicular to clock lines. Each bit occupies 6.24 mil^2 and requires eight IGFETs and two storage capacitors. The entire array including clock generators consumes less than 50 mW of power (nominal).

The functional test sequences contain a total of 3600 vectors and provide 95-percent fault coverage.

The operating range of the chip on the V_{GG} vs V_{SS} plane is shown in Fig. 7. The test was performed at room temperature and nominal clock frequency. The allowed power supply variations are indicated by the box on the plot.

IV. PERFORMANCE

The signal inputs and outputs are all 8-bit μ -law (or, selectably, A-law). The μ -law quantization on these signals limits the echo return loss attainable. The numerical representations of data signals internal to the chip were chosen so as to not significantly degrade the echo-return-loss performance of the chip from that dictated by the μ -law I/O.

Figure 8 shows rate of convergence with white Gaussian noise inputs $x(k)$ having powers of -10 , -20 , and -30 dBmC0 . The procedure for obtaining the plotted measurements is easily explained with the aid of




Fig. 7—Operating range of the echo canceler chip. Boxed-in area indicates allowed variations in system power supplies. Data are taken at room temperature and nominal clock frequency.




Fig. 8—Rate of convergence with a white Gaussian input.

Fig. 9. The reset control is normally at logical high. After adjusting the two variable losses for the desired input signal power and hybrid loss, the reset control is momentarily activated to reset the impulse response estimate to zero. The inhibit control is then released for 100 ms, during which time the echo canceler partially converges. The reading of the

power meter is the datum for echo power after 100 ms. The data for 200, 300, etc., ms are obtained by successively releasing the inhibit control for 100-ms intervals and noting power meter readings.

The simulated echo path in Fig. 9 appears to be a flat-loss path with no frequency dispersion. There is frequency dispersion in the echo path of the actual experiment, however, because of the antialiasing and reconstruction filtering required for the A/D and D/A conversion not shown explicitly.

Since the echo return loss after a set period of adaptation is a random variable, the measurements recorded in Fig. 8 cannot be repeated exactly. Our experience suggests that the standard deviation of the measurements is about 1 dB.

We noted previously that the normalized loop gain of the canceler is such that a 70-dB/s convergence rate is expected. A reference line with this slope is also shown in Fig. 8. The agreement is good.

Table I gives return noise power, echo return loss, and echo return loss enhancement for various values of input signal power and hybrid loss. The method of measurement is the same as that in Fig. 8, but with the reset and inhibit controls held logically inactive at all times. Notice that the entries in Table IC for echo return loss enhancement are nearly constant. The two low values in the lower right can be attributed to the idle-channel-noise floor for 8-bit μ -law encoding. Thus, with the obvious exception that the canceler cannot reduce the noise in the return error signal below the idle-channel-noise floor, the performance of the canceler is well modeled as providing a constant echo return loss enhancement of about 27 dB as the hybrid loss and the input power vary.

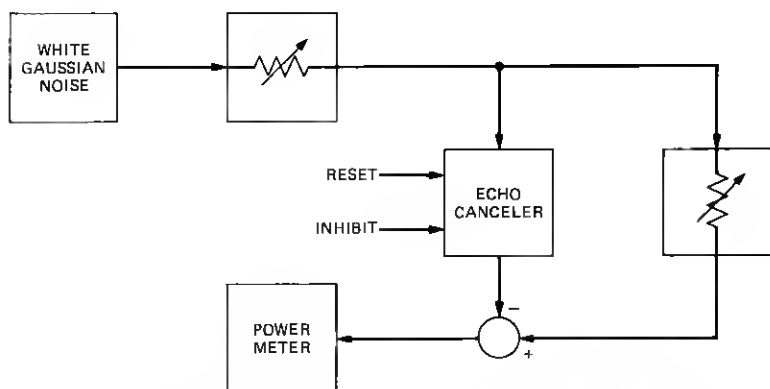


Fig. 9—Procedure to make rate-of-convergence measurements.

Table I—Error signal power, echo return loss, and echo return loss enhancement

(A) Error signal power in dBmC0 as a function of the power in the input signal X in dBmC0 and hybrid loss H in dB

X	H				
	6	9	12	15	18
-10	46	45	41	40	36
-20	36	35	29	29	27
-30	26	24	22	21	20

(B) Echo return loss in dB as a function of the power in the input signal X in dBmC0 and hybrid loss H in dB

X	H				
	6	9	12	15	18
-10	34	35	39	40	44
-20	34	35	41	41	43
-30	34	36	38	39	40

(C) Echo return loss enhancement in dB as a function of the power in the input signal X in dBmC0 and hybrid loss H in dB

X	H				
	6	9	12	15	18
-10	28	26	27	25	26
-20	28	26	29	26	25
-30	28	27	26	24	22

V. ACKNOWLEDGMENTS

Numerous people contributed to the development of the VLSI echo canceler. The initial logic design was in TTL. Y. G. Tao helped in the translation of this logic to NMOS. Some early simulation of the NMOS logic was done by A. Maione. E. T. Grinthal gave us guidance on testing requirements, and E. Heinlein wrote the programs to control the production testing machine. Two on-chip clock drivers capable of driving 250 pf capacitive loads at 1 MHz are required. W. T. Lynch designed these clock drivers and also advised us on the design of a suitable clock distribution scheme. J. R. Barner had full responsibility for the chip layout, and it is much to his credit that the initial marks produced totally working models. Finally, without the administrative support of D. M. Brady, J. M. Sipress, W. D. Warters, R. M. Jacobs, and especially M. R. Aaron, the chip could never have been realized.

REFERENCES

1. H. R. Huntley, "Transmission Design of Intertoll Telephone Trunks," B.S.T.J., 32 No. 5 (September 1953), pp. 1019-1036.
2. R. W. Hatch and A. E. Ruppel, "New Rules for Echo Suppressors in the DDD Network," Bell Laboratories Record, 52 (December 1974), pp. 351-357.
3. G. K. Helder and P. C. Lopiparo, "Improving Transmission in Domestic Satellite Circuits," Bell Laboratories Record, 46 (September 1977), pp. 202-207.
4. J. L. Kelly, Jr. and B. F. Logan, Jr., "Self-Adaptive Echo Canceler," U.S. Patent No. 3,500,000, March 10, 1970.
5. M. M. Sondhi, "Closed Loop Adaptive Echo Canceler Using Generalized Filter Networks," U.S. Patent No. 3,499,999, March 10, 1970.
6. S. J. Campanella, H. G. Suyderhoud, and M. Onufry, "Analysis of an Adaptive Impulse Response Echo Canceler," Comsat Technical Review, 2 (Spring 1972), pp. 1-36.
7. H. Hoge, "Analysis of an Adaptive Echo Canceler with Optimized Gradient Gain," Siemens Forsch, 4, No. 3 (1975), pp. 127-131.
8. N. Demytko and L. K. Mackechnie, "A High-Speed Adaptive Echo Canceler," Aust. Telecom Res., 7 (1973), p. 20.
9. N. Demytko and K. S. English, "Echo Cancellation on Time-Variant Circuits," Proc. IEEE, 65 (March 1977), pp. 444-453.
10. H. Brueggemann, G. Kidd, L. K. Mackechnie, and A. J. Seyler, "An Experimental Echo Canceler for Long Distance Telephone Circuits," Radio Electron Eng., 43 (April 1973), pp. 267-278.
11. Y. Kato et al., "A Digital Adaptive Echo Canceler," NEC Res. Develop., October 1973, No. 31, pp. 32-41.
12. K. Ochiai, T. Araseki, and T. Ogebara, "Echo Canceler with Two Echo Path Models," IEEE Trans. Commun., COM-25 (June 1977), pp. 589-595.
13. S. Tanaka, T. Ogihara, T. Araseki, and K. Ochiai, "A Multichannel Echo Canceler System," NEC Res. Develop., No. 49, (April 1978), pp. 58-64.
14. T. Araseki, Y. Itoh, and K. Ochiai, "A Microprocessor Echo Canceler," Proc. Fourth International Conference on Digital Satellite Communications, Montreal, October, 1978.
15. O. A. Horna, "Echo Canceler with Adaptive Transversal Filter Utilizing Pseudo-Logarithmic Coding," Comsat Technical Review, 7 (Fall 1977), pp. 393-428.
16. D. L. Duttweiler, "A Twelve-Channel Digital Echo Canceler," IEEE Trans. Commun., COM-26 (May, 1978), pp. 647-53.
17. B. Widrow, "Adaptive Filters, I: Fundamentals," Stanford Electronics Lab. Tech. Rept. 6764-6, December, 1966.
18. M. M. Sondhi, "An Adaptive Echo Canceler," B.S.T.J., 46, No. 3 (March 1967), pp. 497-511.
19. R. D. Gitlin and J. S. Thompson, "A Phase-Adaptive Structure for Echo Cancellation," IEEE Trans. Commun., COM-26 (August 1978), pp. 1211-1220.